CRD
Computational Research Division Report
APRIL 2008

Code Booster

Code optimization research explores multicore computing, wins award

A research paper exploring ways to make a popular scientific analysis code run smoothly on different types of multi-core computers won a Best Paper Award at the IEEE International Parallel and Distributed Processing Symposium (IPDPS) this month.

The paper's lead author and CRD researcher, Samuel Williams, and his collaborators chose the lattice Bolzmann code to explore a broader issue: how to make best use of multicore supercomputers. The multicore trend started recently, and the computing industry is expected to

continued on page 3

(a)

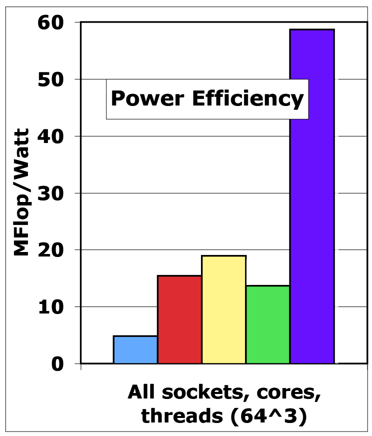

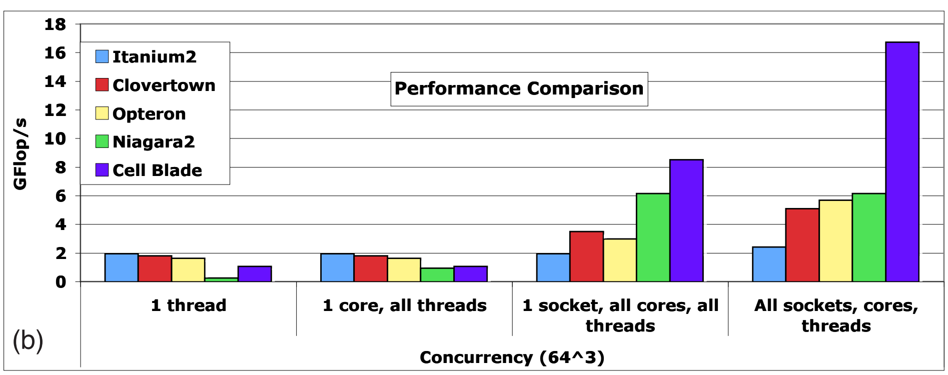

(b)

Comparison of (a) power efficiency and (b) runtime performance across all studied architectures for the 643 problem.

Flood Zone

Expert contributes to two national reports on impact of climate change on transportation

Climate change is expected to lead to flooding of roads, railways, airport runways and other transit systems in the United States, according to two recent national reports that were co-authored by CRD climate modeling expert Michael Wehner.

Though the impacts of climate change will vary by region throughout the United

continued on page 5

Radical Computer

Researchers propose a new breed of supercomputers, aim to improve global climate predictions

Three CRD researchers have proposed an innovative way to improve global climate change predictions by using a supercomputer with low-power embedded microprocessors, an approach that would overcome limitations posed by today's conventional supercomputers.

In a paper published in the May issue of the International Journal of High Performance Computing Applications, Michael Wehner, Lenny Oliker and John Shalf lay out the benefit of a new class of supercomputers for modeling climate conditions and understanding climate change.

continued on page 2

Promoting Diversity

Computing Sciences researchers set out to attract woman and minority candidates

Computing Sciences researchers are undertaking initiatives to work on increasing diversity and to encourage K-12 students to study and pursue a career in math and science.

They formed the Computing Sciences Diversity Committee to provide a forum for researchers and other staff to brainstorm and carry out ideas that promote

continued on page 4

Climate Computer

continued from page 1

Using the embedded microprocessor technology used in cell phones, iPods, toaster ovens and most other modern day electronic conveniences, they propose designing a cost-effective machine for running these models and improving climate predictions.

Understanding how human activity is changing global climate is one of the great scientific challenges of our time. Scientists have tackled this issue by developing climate models that use the historical data of causes and results of the earth's climate, such as rainfall, hurricanes, sea surface temperatures and carbon dioxide in the atmosphere. One of the greatest challenges in creating these models, however, is to develop accurate cloud simulations.

Although cloud systems have been included in climate models in the past, they lack the details that could improve the accuracy of climate predictions. In their research, Wehner, Oliker and Shalf set out to establish a practical estimate for building a supercomputer capable of creating climate models at 1-kilometer scale. A cloud system model at the 1-km scale would provide rich details that are not available from existing models.

To develop a 1-km cloud model, scientists would need a supercomputer that is 1,000 times more powerful than what is available today, the researchers say. But building a supercomputer powerful enough to tackle this problem is a huge challenge.

Historically, supercomputer makers build larger and more powerful systems by increasing the number of conventional microprocessors — usually the same kinds of microprocessors used to build personal computers. Although this approach is feasible for building computers large enough to solve many scientific problems, a system capable of modeling cloud systems at a 1-km scale would cost about $1 billion using the same approach. The system also would require 200 megawatts of electricity to operate, enough energy to power a small city of 100,000 residents.

In their paper, "Towards Ultra-High Resolution Models of Climate and Weather," the researchers present a radical alternative that would cost less to build and require less electricity to operate. They conclude that a supercomputer using about 20 million embedded microprocessors would deliver the results and cost $75 million to construct. This "climate computer" would consume less than 4 megawatts of power and achieve a peak performance of 200 petaflops.

"Without such a paradigm shift, power will ultimately limit the scale and performance of future supercomputing systems, and therefore fail to meet the demanding computational needs of important scientific challenges like the climate modeling," Shalf said.

The researchers arrive at their findings by extrapolating performance data from the Community Atmospheric Model (CAM). CAM, developed at the National Center for Atmospheric Research in Boulder, Colorado, is a series of global atmosphere models commonly used by weather and climate researchers.

The "climate computer" is not merely a concept. Wehner, Oliker and Shalf, along with researchers from UC Berkeley, are working with scientists from Colorado State University to build a prototype system in order to run a new global atmospheric model developed at Colorado State.

"What we have demonstrated is that in the exascale computing regime, it makes more sense to target machine design for specific applications," Wehner said. "It will be impractical from a cost and power perspective to build general-purpose machines like today's supercomputers."

SHANGHAI PROFESSORS VISIT LAB

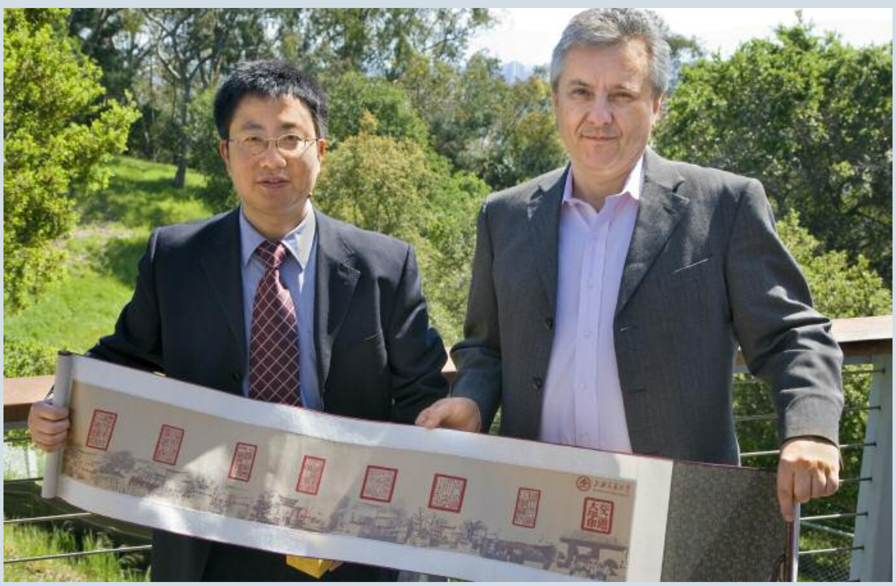

Xiaokang Yang (left), professor and Deputy Director of the Institute of Image Communication and Information Processing at Shanghai Jiao Tong University, presented Horst Simon, CRD Director, with a scroll chronicling the history of the university.

A delegation from Shanghai Jiao Tong University visited Berkeley Lab this month to learn about computational research here. Horst Simon, Associate Lab Director for Computing Sciences and CRD Director, welcomed the group and gave an overview of the CRD research activities.

Shanghai Jiao Tong is one of the premier universities in China, known for its science and engineering programs. The visiting group consisted of professors in electrical engineering and included Professor Wenjun Zhang, who also is vice president of the university.

The group heard presentations on scientific data management by Arie Shoshani, image processing for cryo-electron microscopy by Chao Yang, visualization and analytics by Wes Bethel, machine learning and pattern recognition by Daniela Ushizima, and numerical methods for imaging by James Sethian.

Code Booster *continued from page 1*

add more cores per chip to boost performance in the future. The paper described how the researchers developed a code generator that could efficiently and productively optimize a lattice Bolzmann code to deliver better performance on a new breed of supercomputers built with multicore processors.

The multicore trend is taking flight without an equally concerted effort by software developers. "The computing revolution towards massive on-chip parallelism is moving forward with relatively little concrete evidence on how to best use these technologies for real applications," Williams wrote in the paper.

The researchers settled on the lattice Bolzmann code used to model turbulence in magnetohydrodynamics simulations that play a key role in areas of physics research, from star formation to magnetic fusion devices. The code, LBMHD, typically performs poorly on traditional multicore machines.

The optimization research resulted in a great improvement to the code performance — substantially higher than any published to date. The researchers also gained insight into building effective multicore applications, compilers and other tools.

The paper, "Lattice Boltzmann Simulation Optimization on Leading Multicore Platforms," won the Best Paper Award in the application track. Jonathan Carter of NERSC, Lenny Oliker of CRD, John Shalf of NERSC and Kathy Yelick of NERSC co-authored the paper. The researchers presented their paper at the IPDPS in Miami. Yelick, who is NERSC Director, also was a keynote speaker at the symposium.

Oliker, Carter and Shalf also authored a paper that won the same award last year. The paper, "Scientific Application Performance on Candidate PetaScale Platforms," was co-authored by CRD researchers Andrew Canning, Costin Iancu, Michael Lijewski, Shoaib Kamil, Hongzhang Shan and Erich Strohmaier. Stephane Ethier from the Princeton Plasma Physics Laboratory and Tom Goodale from Louisiana State University also contributed to the work.

In their more recent work on LBMHD, the researchers determined how well the code runs on processors used to build computers today: Intel's quad-core Clovertown, Advanced Micro Devices' dual-core Opteron X2, Sun Microsystems' eight-core Niagra2, and the eight-core STI Cell blade (designed by Sony, Toshiba and IBM). They also looked at Intel's single-core Itanium2 to compare its more complex single core design with other simpler multicores.

The researchers first looked at why the original LBMHD performs poorly on these multicore systems. Williams and his fellow researchers found that, contrary to conventional wisdom, memory bus bandwidth didn't present the biggest obstacle. Instead, lack of resources for mapping virtual memory pages, insufficient cache bandwidth, high memory latency, and/or poor functional unit scheduling did more to hamper the code's performance, Williams said.

The researchers created a code generator abstraction for LBMHD in order to optimize it for different multicore architectures. The optimization efforts included loop restructuring, code reordering, software prefetching, and explicit SIMDization. The researchers characterized their effort as akin to the "auto-tuning methodology exemplified by libraries like ATLAS and OSKI."

The results showed a wide range of performance results on different processors and pointed to bottlenecks in the hardware that prevented the code from running well. The optimization efforts also resulted in a huge gain in performance — the optimized code ran up to 14 times faster than the original version. It also achieved sustained performance for this code that is higher than any published to date: over 50 percent of peak flops on two of the processor architectures.

Compared with other processors, the Cell processor provided the highest raw performance and power efficiency for LBMHD. The processor's design calls for a direct software control of the data movement between on-chip and main memory, resulting in the impressive performance. Overall, the researchers concluded, processor designs that focused on high throughput using sustainable memory bandwidth and a large number of simple cores perform better than processors with complex cores that emphasized sequential performance.

They also concluded that auto-tuning would be an important tool for ensuring that numerical simulation codes will perform well on future multicore computers.

Read about the researchers' analyses of other processor architectures by checking out the paper on Williams' website.

Hall of Fame

Parallel Computing Conference

The 13th SIAM Conference on Parallel Processing for Scientific Computing again drew a large number of CRD researchers, who presented papers on subjects such as hardware and applications for petascale computing, adaptive mesh refinement algorithms, data-flow programming techniques, and power-efficient hardware and software designs.

The three-day conference by the Society for Industrial and Applied Mathematics drew hundreds of attendees from universities, national labs and other research institutions around the world. Esmond Ng, head of CRD's Scientific Computing Group, served as a co-chair for the conference in Atlanta, Georgia last month.

The conference emphasized the intersection between high performance scientific computing and scalable algorithms, architectures and software. It drew researchers from the applied mathematics, computer science, computational science and engineering fields.

More than a dozen CRD scientists contributed to the conference as authors/speakers or organizers of minisymposia. The CRD researchers who participated were John Bell, Andrew Canning, Phillip Colella, James

continued on page 6

Diversity Committee *continued from page 1*

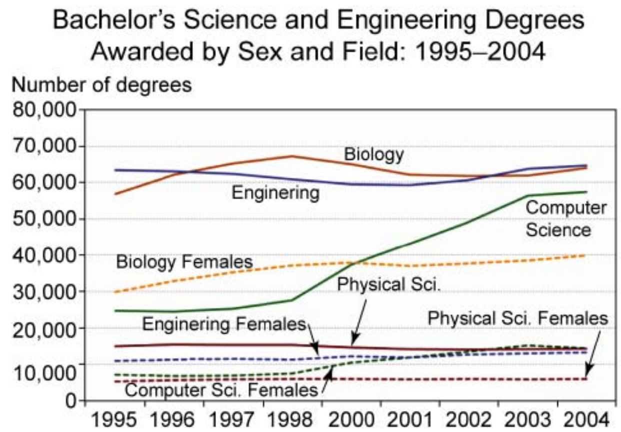

Source: NSF Report nsf07315—Women, Minorities, and Disabilities in Science and Engineering, 2007

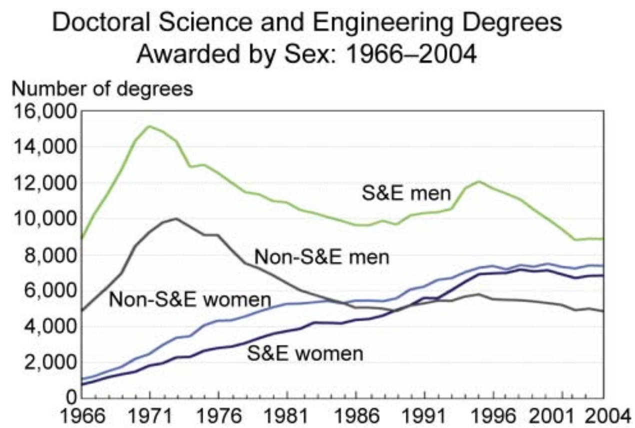

Source: NSF Report nsf07315—Women, Minorities, and Disabilities in Science and Engineering, 2007

diversity. The committee will also coordinate the various diversity-related efforts that many researchers and managers already have undertaken.

Members of the diversity committee include representatives from the groups and departments within the Computational Research Division and NERSC, the two divisions within Computing Sciences (CS).

Computing Sciences staff were actively involved in diversity-related recruitment and programs before the diversity committee was formed late last year. They have attended conferences such as the Grace Hopper Celebration of Women in Computing and the Richard Tapia Celebration of Diversity in Computing. They then followed up by sending student attendees information about research opportunities at Berkeley Lab. Juan Meza, head of the High Performance Computing Research Department, led the Mathematical Sciences Research Institute's 2007 summer program for undergraduate students from underrepresented groups. Three researchers, Cecilia Aragon, Ann Almgren, and Deb Agarwal spoke at the 2007 University of San Francisco Summer Enrichment Program, which set out to encourage young women ages 13–18 to study computer science.

"We have always cared about diversity, and this committee offers an opportunity for all interested CS employees to get involved and give their input," said Agarwal a member of the diversity working group.

The diversity committee plans to carry out projects that would make a difference not only in hiring and retaining women and minorities, but also in making science a fun and rewarding subject in schools.

"Our participation in diversity work groups promotes accessibility and exposure to computer science careers for people who may not otherwise think about pursuing them," said Marcia Ocon, a senior human resources officer for Computing Sciences and a member of the diversity committee. "It's our hope that in the long term, our recruiting efforts will be met with a more diverse applicant pool."

Research has shown that disparity begins in schools, where female students are less likely to pursue a degree in science and engineering. In fact, women were 28.5 percent of the people in the United States earning bachelor's degrees in computer science in 1995, but that figure declined to 25 percent in 2004, according to 2007 National Science Foundation report on "Women, Minorities, Persons with Disabilities in Science and Engineering."

The same report showed that women received 17.2 percent of the Ph.D.s in computer sciences in 1998 and 19.8 percent in 2005. Overall, women comprised 34.3 percent of the earned Ph.D.s in science and engineering in 1998 and 37.7 percent in 2005. The science and engineering category includes social sciences such as psychology and sociology.

As a result, the number of women who qualify for science and engineering research positions is much smaller. Other factors, such as a lack of support for women with children and biased hiring policies, compound the problem. Statistics show that married men with children under age six are 50 percent more likely to enter a tenure track position than married women with children under six, according to a presentation by Mary Ann Mason, Dean of the Graduate Division at UC Berkeley, during a 2007 conference by the American Physical Society. The statistics came from a survey of Ph.D. recipients in sciences and humanities from 1981 to 1995.

Minorities, of course, make up a smaller employee pool in science and engineering. The NSF report showed that among the employed scientists with a Ph.D. in 2003, 25.2 percent of them were minorities. The figure for employed engineers with a Ph.D. was 37.2 percent.

Agarwal noted that women and minorities might not be choosing computer sciences because they believe job opportunities in the field are going overseas. In fact, in India, where she gave a recent presentation on gender and employment, women show a strong interest in computer sciences because of the country's booming software industry.

Members of the diversity committee are eager to get started on several projects, including an internal mentoring program and a "career in computing camp" for local high school students. They also plan to continue to attend diversity conferences and participate in job fairs in order to attract more women and minority candidates.

Check out the National Science Foundation web site that provides a wealth of data on education and employment for women and minorities in science and engineering.

Climate and Floods *continued from page 1*

States, it is certain they will be widespread and costly in human and economic terms. It also will require significant changes in the planning, design, construction, operation, and maintenance of transportation systems, according to the first report, "The Potential Impacts of Climate Change on U.S. Transportation," published by the National Research Council. The second report, "Climate Change and Variability on Transportation Systems and Infrastructure: Gulf Coast Study, Phase I," issued by the U.S. Department of Transportation, focuses on the impact on 48 contiguous counties in four states, from Galveston, Texas to Mobile, Alabama.

"Global warming is very likely to have serious impacts on many aspects of society, some with high economic costs. The transportation sector is just one of these," Wehner said. "Our reports are intended to give government planners the information they need to make the difficult decisions which will be necessary to adapt to these almost certain impacts."

Throughout the country, transportation planners and engineers use historical temperature and precipitation data to help them design transportation systems that can withstand local weather and climate conditions. Both reports find that those climate predictions may no longer be reliable, however, in the face of new weather and climate extremes. Roads, transit systems and airports that were built based on older data could fail as a result.

The first report looks at the impact across the United States and its territories and finds that California airports in San Francisco, Santa Barbara, and Oakland could be inundated under conditions of extreme high tides coupled with flood conditions and exacerbated by local sea-level rise. In the Bay Area, major highways and railroads near sea level could also be threatened by rises in sea level. Not all climate changes will be negative, however. Marine transportation could benefit from more open seas in the Arctic, creating new and shorter shipping routes and reducing transport time and costs. In cold regions, rising temperatures could reduce the costs of snow and ice control and would make travel conditions safer for passenger vehicles and freight.

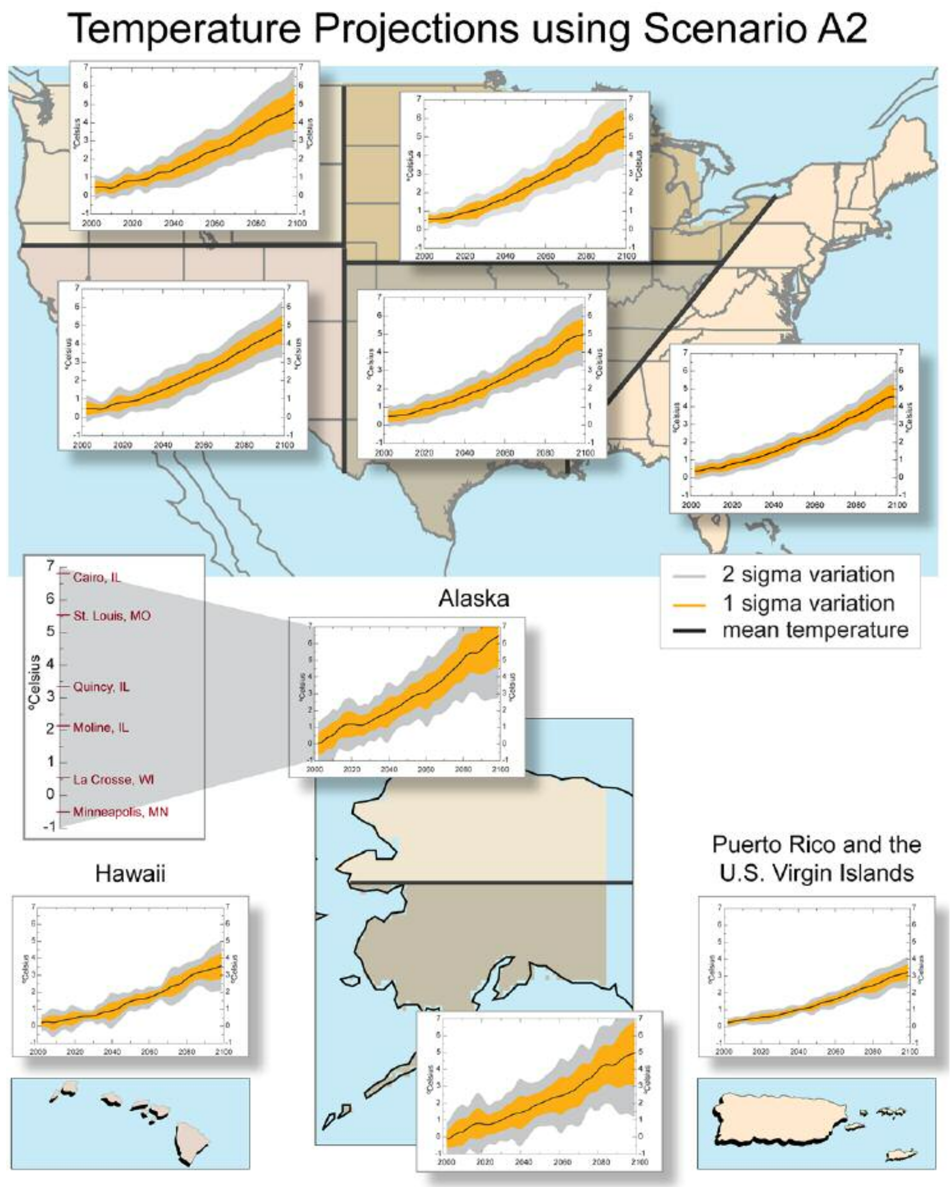

Climate model projections of surface air temperature (in Centigrade) and their associated uncertainties for a business as usual future scenario with little reduction of human-produced greenhouse gas emissions.

The report draws upon five papers commissioned by the Transportation Research Board, which, like the National Research Council, is part of the National Academies. One of the papers, "Climate Variability and Change with Implications for Transportation," was co-authored by Thomas C. Peterson, Marjorie McGuirk and Tamara G. Houston of the National Oceanographic and Atmospheric Administration's (NOAA) National Climatic Data Center, Andrew H. Horvitz of NOAA's National Weather Service and Wehner.

The second report, focusing on the Gulf Coast, provides an assessment of the vulnerabilities of transportation systems in the region to potential changes in weather patterns and related impacts, as well as the effect of natural land subsidence and other environmental factors

continued on page 6

Climate and Floods

continued from page 5

in the region. The authors of the study find that changes in climate could disrupt transportation services in the region. Twenty-seven percent of major roads, 9 percent of rail lines and 72 percent of area ports could be vulnerable to flooding due to future sea level rise.

For both reports, Wehner and his co-authors first analyzed observed climate data (daily temperature and precipitation) from 1950 to 2005. Then they used the large collection of climate model data, archived at the Program for Climate Model Diagnosis and Intercomparison (PCMDI) at Lawrence Livermore National Laboratory, to validate the models against observations and to create plausible projections of future climate change relevant to transportation analysts. They considered future scenarios ranging from very aggressive reduction of human-produced greenhouse gas emissions to none at all.

The simulation data, based on 25 climate models, reflect the climate research community's collective efforts to predict the global climate of the future. The simulations took into account different factors based on economical, technological, and sociological assumptions. Researchers from the U.S., Norway, Canada, France, Australia, Russia, Germany, Korea, China, Japan and the United Kingdom contributed to the collection at PCMDI. The Intergovernmental Panel on Climate Change (IPCC) used these data for its Fourth Assessment Report released last year, a comprehensive document that won the IPCC a share of the 2007 Nobel Peace Prize.

Hall of Fame

continued from page 1

Demmel, Tony Drummond, Parry Husbands, Sherry Li, Esmond Ng, Lenny Oliker, John Shalf, Horst Simon, Philip Sternberg, Brian Van Straalen, and Chao Yang.

One of the highlights at the conference was a panel discussion organized by Horst Simon. The intent of the panel was to provide a global perspective on parallel processing research for scientific computing. There were representatives from Brazil, China, India, Germany, Japan, South Africa, and the U.S.A. The panelists provided overviews of high-performance scientific computing in various countries and touched on issues related to opportunities for global collaborations.

Greening of HPC

Horst Simon, Associate Lab Director for Computing Sciences at Berkeley Lab, recently gave a talk as part of the Distinguished Lecture Series in Petascale Simulation at the University of Texas at Austin.

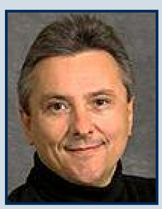

Horst Simon

Simon discussed efforts to promote the growth of high performance computing without contributing to global warming in the talk, "The Greening of High Performance Computing — Will Power Consumption Become the Limiting Factor for Future Growth?" He also outlined the Lab's research projects that address the issue of reducing power consumption.

ACM Award

Dave Patterson, a CRD computer scientist and a UC Berkeley professor, won the 2007 ACM Distinguished Service Award.

Dave Patterson

The Association of Computing Machinery gave Patterson the award for his initiatives that brought more respect and understanding to the computing profession.

Patterson is the founding director of the Parallel Computing Laboratory (Par Lab) at UC Berkeley. Par Lab works on projects that aim to solve hardware and software challenges in multicore computing. He also founded the Reliable, Adaptive and Distributed Systems Laboratory (RAD Lab), which focuses on the design of more dependable computing systems.

His membership on President's Information Technology Advisory Committee (PITAC) under President George W. Bush and on the National Academy of Engineering (NAE) Membership Committee enabled him to advance the interest of the computing research community. He helped to increase the number of computer scientists being elected to the NAE annually by highlighting work by outstanding computing researchers.

About CRD Report

CRD Report, which publishes every other month, highlights the cutting-edge research conducted by staff scientists in areas including turbulent combustion, nanomaterials, climate change, distributed computing, high-speed networks, astrophysics, biological data management and visualization. CRD Report Editor Ucilia Wang can be reached at 510 495-2402 or **Uwang@lbl.gov**. Find previous CRD Report articles at http://crd.lbl.gov/html/news/CRDreport.html.